

At the outset, Applicants note with appreciation the indication that claims 3, 7, and 14 would be allowable if rewritten in independent form to include the features of the base claim and any intervening claims.

Applicants propose amending Figs. 3A and 3B to correct typographical errors. The proposed changes are marked in red in the attached Request for Approval of Drawing Change. Applicants respectfully request that the Examiner approve the corrections to Figs. 3A and 3B.

In the Office Action, the Examiner requests that the Information Disclosure Statement, filed October 18, 1999, be resubmitted. In response, Applicants resubmit the Information Disclosure Statement and respectfully request that the Examiner consider the documents cited therein.

The Examiner rejects claims 22 and 23 under 35 U.S.C. § 112, second paragraph, as allegedly indefinite for failing to particularly point out and distinctly claim the subject matter which Applicants regard as the invention. More specifically, the Examiner alleges that the claims are generally narrative and indefinite and fail to conform to current U.S. practice (Office Action, pg. 2). In an attempt to expedite prosecution, Applicants cancel claims 22 and 23 herewith, thereby rendering the rejection of these claims moot.

The Examiner rejects claims 1, 2, 4-6, 8-13, and 15-23 under 35 U.S.C. § 102(e) as allegedly anticipated by DOSHI et al. Applicants respectfully traverse this rejection.

DOSHI et al. is directed to a distributed precomputation algorithm for determining primary and/or restoration paths in an optical or electrical network (Abstract). In DOSHI et al., the precomputation algorithm provides end-to-end path restoration (col. 10, lines 36-37). The end nodes detect failures in a path (col. 10, lines

37-43) and an alternate route using the restoration path can be triggered without trying to evaluate exactly where the failure has occurred (col. 10, line 67, to col. 11, line 3).

In contrast, Applicants' invention recited, for example, in independent claim 1 recites that at least one of the nodes in the network comprises a processor to compute an alternative route for an initial route, a storage space to store the initial route and the alternative route, a mechanism to detect failure in a downstream network element in the initial route, and a forwarder to automatically forward a packet to the next node. DOSHI et al. does not disclose or suggest this combination of features.

For example, DOSHI et al. does not disclose or suggest a mechanism to detect failure in a downstream network element in the initial route. The Examiner relied on col. 10, lines 53-55, of DOSHI et al. for allegedly disclosing this feature (Office Action, pg. 3). Applicants submit that this section of DOSHI et al. does not disclose or suggest the recited feature.

Col. 10, lines 53-55, of DOSHI et al. discloses "[a]s noted above, failures are generally detected in all-optical networks by measurements of optical signal power levels." This section of DOSHI et al. merely discloses how failures in an all-optical network are detected. This section of DOSHI et al. does not disclose or suggest, however, that the detection of failures are performed by a mechanism within at least one node in the network, as required by Applicants' claim 1.

In fact, DOSHI et al. discloses that detection of failures is performed by the end nodes (i.e., the source and destination devices) – see, for example, col. 10, lines 40-41. Therefore, the intermediate network nodes in DOSHI et al. (i.e., those nodes located between the source and destination nodes) would not include a mechanism to detect

failure in a downstream network element in the initial route, as recited in Applicants' claim 1, since DOSHI et al. specifically discloses that detection of failures in a given path is performed by the source and destination devices. Moreover, DOSHI et al. discloses at col. 10, line 67, to col. 11, line 3, that an alternate route using the restoration path can be triggered without trying to evaluate exactly where the failure occurred. This section of DOSHI et al. teaches away from incorporating the recited mechanism to detect failure into network nodes of the DOSHI et al. system.

A proper rejection under 35 U.S.C. § 102 requires that the reference teach every aspect of the claimed invention either explicitly or impliedly. Any feature not directly taught must be inherently present. See M.P.E.P. 2131. Since, as set forth above, DOSHI et al. does not disclose at least one node in the network comprising a mechanism to detect failure in a downstream network element in the initial route, Applicants respectfully submit that the rejection of claim 1 under 35 U.S.C. § 102(e) as anticipated by DOSHI et al. should be withdrawn.

For at least the foregoing reasons, Applicants submit that DOSHI et al. does not anticipate claim 1.

Claims 2 and 4-6 depend from claim 1. Accordingly, Applicants submit that these claims are not anticipated by DOSHI et al. for at least the reasons given above with respect to claim 1. Moreover, these claims recite additional features not disclosed or suggested by DOSHI et al.

For example, claim 6 recites that the mechanism to detect failures sends communication packets to downstream nodes at regular intervals. The Examiner relied on col. 14, lines 50-52, of DOSHI et al. for allegedly disclosing this feature (Office

Action, pg. 4). Applicants submit that this section of DOSHI et al. does not disclose or suggest the feature of claim 6.

Col. 14, lines 50-52, of DOSHI et al. discloses "[t]he algorithm may be restarted from scratch after major provisioning or network capacity changes, or on a periodic, but infrequent basis." This section of DOSHI et al. teaches away from sending communication packets to downstream nodes at regular intervals by disclosing that the precomputation algorithm can be restarted on a periodic, but infrequent basis.

For at least this additional reason, Applicants submit that claim 6 is not anticipated by DOSHI et al.

Independent claim 8 recites, *inter alia*, detecting a failed element and automatically forwarding packets on the alternative route without communicating with either the source or the destination. The Examiner grouped the features of claim 8 with those of independent claim 1 (Office Action, pg. 3). As set forth above with respect to claim 1, DOSHI et al. discloses that detection of failures is performed by the end nodes (i.e., the source and destination devices) – see, for example, col. 10, lines 40-41. DOSHI et al. also discloses that the end nodes initiate the path restoration – see, for example, col. 3, lines 42-44, that discloses "[i]n end-to-end path based restoration, the source-destination node pairs of demands affected by a given failure initiate the restoration action." Therefore, DOSHI et al. does not disclose automatically forwarding packets on the alternative route without communicating with either the source or the destination since DOSHI et al. discloses that the source and destination devices are involved in the detection of failures and the initiation of path restoration.

For at least the foregoing reason and for those reasons given above with respect to claim 1, Applicants submit that DOSHI et al. does not anticipate claim 8.

Claims 9-13 and 15-17 depend from claim 8. Therefore, Applicants submit that these claims are not anticipated by DOSHI et al. for at least the reasons given above with respect to claim 8. Moreover, these claims recite additional features not disclosed or suggested by DOSHI et al.

For example, claim 12 recites that that the detecting a failure is conducted locally by a node preceding the failed element without requiring notification of a master server or an ingress node. The Examiner relied on col. 10, lines 37-43, of DOSHI et al. for allegedly disclosing this feature. Applicants submit that this section of DOSHI et al. does not disclose or suggest the feature of claim 12.

Col. 10, lines 37-43, of DOSHI et al. discloses:

The illustrative algorithm is distributed in nature in that restoration paths are computed and stored locally at each node, restoration path routing tables are activated locally at each node, failure is detected on a path by two end nodes independently, and the path restoration does not require intervention of a central controller.

This section of DOSHI et al. discloses that failure in a path is detected by the two end nodes independently, and the path restoration does not require intervention of a central controller. DOSHI et al.'s end nodes are not, however, "a node preceding the failed element," as recited in Applicants' claim 12. This section of DOSHI et al. teaches away from the feature recited in claim 12 by disclosing that the detection of a failure is performed by the end nodes.

For at least this additional reason, Applicants submit that claim 12 is not anticipated by DOSHI et al.

Independent claim 18 recites features similar to those recited above with respect to claim 1. Therefore, Applicants submit that claim 18 is not anticipated by DOSHI et al. for reasons similar to those given above with respect to claim 1.

Claims 19-21 depend from claim 18. Therefore, Applicants submit that these claims are not anticipated by DOSHI et al. for at least the reasons given above with respect to claim 18.

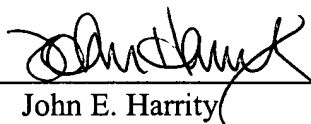
Claims 22 and 23 have been canceled herewith thereby rendering the rejection of those claims moot.

In view of the foregoing amendment and remarks, Applicants respectfully request the Examiner's reconsideration of this application, and the timely allowance of the pending claims.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectively submitted,

HARRITY & SNYDER, L.L.P.

By: 
John E. Harrity
Registration No. 43,367

Date: December 27, 2002

11240 Waples Mill Road
Suite 300
Fairfax, Virginia 22030
(571) 432-0800

ATTACHMENT SHOWING CHANGES MADE

IN THE SPECIFICATION:

The paragraph at page 8, line 19, has been amended as follows:

Referring back to Figure 3a, after initialization, alternative routes leading from the alternative-route-enabled nodes are computed (308). The alternative routes are computed along the initial route computed in step 330 of Fig. 3b to the destination device 110. Information associated with the alternative route is stored locally at the alternative-route-enabled nodes so an alternative route is readily available for rapid resumption of network traffic when a failure occurs in the initial route. Packets are forwarded using the initial routes (310). According to the stored topology information, each node forwards packets to the IP address of the next node on the specified link. In a MPLS framework, each node swaps a label on the IP packet before sending it to the next hop. While in session, RSVP hello extensions are implemented at each node providing hello packets that can be sent directly to other nodes on the initial route. A failure can thus be quickly detected (by a failure to respond to a hello message). Once a failure is detected (315), and if the detecting node is alternative-route-enabled (322), traffic for the failed route is directed to the pre-computed alternative route stored locally (320). Thereafter, packets are forwarded using the alternative route. If the detecting node is not alternative-route-enabled, a message is sent to an ingress node (324), the first node in the system on the initial route. Thereafter, the process ends ([330]326) and the ingress node may invoke a back-up mechanism. In another implementation, when the node detecting a failure is not alternative-route-enabled, a failure message can be forwarded upstream to a nearest in-

line alternative-route-enabled node on the initial route. The alternative-route enabled node can implement the alternative route and continue forwarding packets.

The paragraph at page 9, line 10, has been amended as follows:

If no failure is detected, packets are forwarded along the initial route or alternative route until the traffic ends at step 325 and the process finishes in step [330]326. Single or multiple failures in the system can be supported by directing traffic to alternative routes stored locally at alternative-route-enabled nodes.

IN THE CLAIMS:

Claims 1, 4, and 18 have been amended as follows:

1. (Amended) A network for forwarding packets from a source device to a destination device, said network including a plurality of network elements including nodes and connecting links, a master server for monitoring the network and establishing an initial route between the source device and the destination device, wherein at least one [node] of the nodes comprises:

a processor to compute [an] at least one alternative route for the initial route;

a storage space to store the initial route and the at least one alternative route;

a mechanism to detect failure in a downstream network element in the initial route; and

a forwarder to automatically forward a packet to the next node on one of
the at least one alternative route.

4. (Amended) The node of claim 1, wherein the processor computes an
alternative route not including the downstream [node and link on] network element in the
initial route.

18. (Amended) A method for locally rerouting packets traveling on an
established route when a node in a network of interconnected nodes fails, the method
comprising:

computing, at a plurality of intermediary nodes along the [initial]
established route, an alternative route leading from the computing node to the destination
device of the established route;

determining locally that the established route has failed; and

automatically forwarding packets [according] on the alternative route.